

Acceleration of matrix element computations for precision measurements[☆]

O. Brandt^{a,b,*}, G. Gutierrez^c, M.H.L.S. Wang^c, Z. Ye^d

^a*II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen, Germany*

^b*now at Kirchhoff-Institut für Physik, Universität Heidelberg, Heidelberg, Germany*

^c*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

^d*University of Illinois at Chicago, Chicago, Illinois 60607, USA*

Abstract

The matrix element technique provides a superior statistical sensitivity for precision measurements of important parameters at hadron colliders, such as the mass of the top quark or the cross section for the production of Higgs bosons. The main practical limitation of the technique is its high computational demand. Using the concrete example of the top quark mass, we present two approaches to reduce the computation time of the technique by a factor of 90. First, we utilize low-discrepancy sequences for numerical Monte Carlo integration in conjunction with a dedicated estimator of numerical uncertainty, a novelty in the context of the matrix element technique. Second, we utilize a new approach that factorizes the overall jet energy scale from the matrix element computation, a novelty in the context of top quark mass measurements. The utilization of low-discrepancy sequences is of particular general interest, as it is universally applicable to Monte Carlo integration, and independent of the computing environment.

Keywords: matrix element, Monte Carlo integration, low-discrepancy sequences, hadron collider, top quark.

PACS numbers: 02.60.Jh, 02.70.Uu, 02.50.Sk, 14.65.Ha.

1. Introduction

The matrix element (ME) technique [1] is a powerful tool in experimental particle physics, especially at hadron colliders, as it provides a superior statistical sensitivity in the extraction of important parameters of the standard model. This sensitivity is achieved by taking into account the full topological and kinematic information in a given event, and determining the probabilities P_{sig} and P_{bkg} for observing each event, assuming respective signal and background hypotheses in the respective ME probabilities $|\mathcal{M}_{\text{sig}}|^2$ and $|\mathcal{M}_{\text{bkg}}|^2$. In the context of searches for new physics, these probabilities can be used to construct the most powerful test statistic $Q \equiv \frac{P_{\text{sig}}}{P_{\text{bkg}}}$ according to the Neyman-Pearson lemma [2]. An advantage of the ME technique is that it calculates P_{sig} and P_{bkg} *ab initio*, in contrast to multivariate methods. Furthermore, P_{sig} depends directly on the physical parameter of interest in a specific theoretical framework.

The ME technique was first suggested by Kondo [1] and pioneered in the context of experimental particle physics at the Tevatron in the measurement of the mass of the top quark m_t [3], in the determination of the helicity of the W boson [4], as well as for the first evidence for production of single top quarks [5, 6]. Since then, the ME technique has been used in several analyses, for example in searches for the Higgs boson at the Tevatron [7] and at the LHC [8]. Recently, a general framework for the ME technique, named MadWeight [9], has become available.

Despite its superior statistical sensitivity, the ME technique is not widely applied because of its high computational demand. For example, to perform a previous measurement of m_t using 3.6 fb^{-1} of integrated luminosity [10] by the D0 Collaboration, about two million CPU-hours were required on a single core of the 64 bit XEON E5-2620 CPU, with a clock rate of 2 GHz, and a 64 bit computation. In this manuscript, we present two approaches that were successfully applied to reduce the computational demand of the ME technique by two orders of magnitude. First, we utilize low-discrepancy sequences (LDS) for numerical Monte Carlo (MC) integration, in conjunction with a dedicated estimator of the numerical uncertainty, which is a novelty in the context of the ME technique. Second, we factorize the overall jet energy scale (JES) from the ME computation, which was never done before in the context of m_t measurements using an *in situ* JES calibration. The use of LDS is generally applicable to MC integration. In particular, this approach is not hardware-specific, i.e., it can be used on, e.g., a graphical processing unit.

We present our results using the example of the recent measurement of the top quark mass [11], the single most precise measurement of this parameter, yielding $m_t = 174.98 \pm 0.58 \text{ (stat)} \pm 0.49 \text{ (syst)} \text{ GeV}$. This measurement was performed in lepton+jets final states¹ with the full sample of $p\bar{p}$ collision data from the Fermilab Tevatron Collider at $\sqrt{s} = 1.96 \text{ TeV}$, corresponding to 9.7 fb^{-1} of integrated lumi-

[☆]FNAL report number: Fermilab-Pub-14/416-E.

*Corresponding author

Email address: obrandt@fnal.gov (O. Brandt)

¹The lepton+jets final states aim at selecting the $p\bar{p} \rightarrow t\bar{t} \rightarrow W^+bW^-\bar{b} \rightarrow \ell^+ \nu b q \bar{q}' \bar{b}$ and its charge conjugate process, where t and b denote respectively top and bottom quarks, W^\pm is the W boson, ℓ^\pm stands for charged leptons, and ν represents a neutrino.

osity. The computational demand arises not so much from the number of events recorded in $p\bar{p}$ collisions, but rather from number of the simulated MC events which are used for the calibration of the method and for the evaluation of systematic uncertainties. D0's previous measurements of m_t [10] and of the difference $\Delta m = m_t - m_{\bar{t}}$ [12], both using 3.6 fb^{-1} of integrated luminosity, were also performed with the ME technique.

This manuscript is structured as follows. We begin with a brief review of our previous implementation of the ME technique for the measurement of m_t [10] in 3.6 fb^{-1} of data. This analysis applies several approaches to reduce the computational demand that potentially have general interest. We follow with a discussion of our latest implementation of the ME technique, which provides further reduction in the computational demand through use of LDS for the MC integration, presented in Sec. 3, and through factorization of the scale factor for jet energies k_{JES} from the ME computation, discussed in Sec. 4. Finally, we present in Sec. 5 the validation of our latest implementation of the ME technique with pseudo-experiments (PE), comprised of MC events fully simulated in the D0 detector, and conclude in Sec. 6. The MC simulations are described in Ref. [11].

2. Previous implementation of the matrix element technique

The extraction of m_t with the ME technique is performed with a likelihood that uses per-event probability densities (PD) defined by the ME of the processes contributing to the observed events. Assuming two non-interfering contributions from $t\bar{t}$ and $W + \text{jets}$ production, the per-event PD is given by

$$P_{\text{evt}} = A(\vec{x})[fP_{\text{sig}}(\vec{x}; m_t, k_{\text{JES}}) + (1-f)P_{\text{bkg}}(\vec{x}; k_{\text{JES}})], \quad (1)$$

where the observed signal fraction f , m_t , and the overall multiplicative factor k_{JES} adjusting the energies of jets after their default jet energy scale calibration, are parameters to be determined from data. The \vec{x} denotes the measured jet and lepton four-momenta, and $A(\vec{x})$ accounts for acceptance and efficiencies. The function P_{sig} represents the PD for $t\bar{t}$ production, and P_{bkg} refers to the PD for $W + \text{jets}$ production.

In general, the measured set \vec{x} will not be identical to the set of corresponding partonic variables \vec{y} because of finite detector resolution and parton hadronization. Their relationship is described by a transfer function $W(\vec{x}, \vec{y}, k_{\text{JES}})$. The densities P_{sig} and P_{bkg} are calculated through a convolution of the differential partonic cross section, $d\sigma(\vec{y})$, with $W(\vec{x}, \vec{y}, k_{\text{JES}})$ for the final-state partons and the PD for the initial-state partons, $f(q_i)$, where the q_i are the momenta of the colliding partons. This is done by integrating over all possible parton states that lead to \vec{x} :

$$P_{\text{sig}}(\vec{x}; m_t, k_{\text{JES}}) = \frac{1}{\sigma_{t\bar{t}, \text{obs}}(m_t, k_{\text{JES}})} \int \sum d\sigma(\vec{y}, m_t) d\vec{q}_1 d\vec{q}_2 \times f(\vec{q}_1) f(\vec{q}_2) W(\vec{x}, \vec{y}; k_{\text{JES}}). \quad (2)$$

The sum extends over all possible flavor combinations of the initial-state partons. The longitudinal-momentum parton density functions (PDF) $f(q_{i,z})$ are taken from the CTEQ6L1

set [13], while the dependencies $f(q_{i,x})$, $f(q_{i,y})$ on transverse momenta follow those PD obtained from the PYTHIA simulation [14, 15]. The factor $\sigma_{t\bar{t}, \text{obs}}(m_t, k_{\text{JES}})$, defined as the total cross section for $t\bar{t}$ production in $p\bar{p}$ collisions to be observed in the detector, ensures that $A(\vec{x})P_{\text{sig}}$ is normalized to unity. The differential cross section, $d\sigma(\vec{y}, m_t)$, is calculated using the leading order (LO) ME for the process $q\bar{q} \rightarrow t\bar{t}$ [16, 17].

The calculation in Eq. 2 at LO involves 24 integration variables associated with the two initial-state partons and the six partons in the final state. The directions of the four jets and the charged lepton in (η, ϕ) space are well-measured, and are therefore represented by ten δ -functions. After accounting for these δ -functions, and imposing energy-momentum conservation through four additional δ -functions, ten integration variables remain.

The integration in Eq. 2 is performed numerically using the MC integration method of Ref. [18]. The pseudo-random numbers for the MC integration are generated with RANLUX [19] in a $[0, 1]^{10}$ hypercube, and then transformed to the ranges of the integration variables. Importance sampling [20] is utilized to reduce the computational demand of the integration. Furthermore, we perform a Jacobian transformation of the nominal ten integration variables to variables where prior information is either known or can be easily obtained. This prior information is then used in importance sampling. The optimized integration variables are: m_{W^+} , m_{W^-} , m_t , $m_{\bar{t}}$, $q_{1,x}$, $q_{1,y}$, $q_{2,x}$, $q_{2,y}$, $\rho = E_q/(E_q + E_{q'})$ for the quarks from $W \rightarrow q\bar{q}'$ decay in the LO picture where E represents the particle's energy, and, the energy (curvature) of the electron (muon track) κ .

To integrate over m_t and $m_{\bar{t}}$, random numbers are generated according to expected Breit-Wigner distributions for each given m_t hypothesis. The constraint of $M_W = 80.4 \text{ GeV}$ for the *in-situ* JES calibration is imposed by integrating over W boson masses using a Breit-Wigner prior. For the integration over $q_{i,x}$ and $q_{i,y}$, the ME is sampled in transverse momentum $p_T^{q_i}$ according to the distribution predicted in MC simulations, and uniformly in ϕ^{q_i} . To integrate over κ , random numbers are generated according to the corresponding part of the transfer function, which is defined as the probability to obtain the measured κ_x value, given a value κ_y at the parton level.

Importance sampling in ten bins is employed for the integration over ρ . The MC integration is performed iteratively with an increasing number of samplings of the integral per iteration, where each iteration uses the probability distribution in ρ from the previous one as input for importance sampling.

There are 24 possible jet-parton assignments that are summed with weights based on their consistency with b -tagging information². Typically, two and sometimes four or six jet-parton assignments numerically dominate the final result for P_{sig} . To identify them, we perform a pre-integration step, where we calculate P_{sig}^i for each jet-parton assignment i , until a relative numerical precision of 10% is reached, or the integral is sampled $2^{14} = 16,384$ times. The numerical precision of those jet-parton assignments with P_{sig}^i within 2% of the maximal P_{sig}^i

²We identify jets from b quarks through the use of a multivariate algorithm, as discussed in Ref. [11].

value is further refined until the desired precision has been achieved, or the integral is sampled $2^{24} = 16,777,216$ times. For all other assignments P_{sig}^i obtained in the pre-integration step is kept.

The differential partonic cross section for P_{bkg} is calculated similarly to P_{sig} , i.e., applying MC integration and the same transfer function $W(\vec{x}, \vec{y}; k_{\text{JES}})$, however using the LO $W + 4$ jets ME implemented in VECBOS [21]. Here, the initial-state partons are all assumed to have no transverse momentum $p_T = 0$.

To extract m_t and k_{JES} , we calculate P_{sig} and P_{bkg} on a grid in (m_t, k_{JES}) with spacings of $(1 \text{ GeV}, 0.01)$. A likelihood function $\mathcal{L}(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N; m_t, k_{\text{JES}}, f)$ is constructed at each grid point from the product of the individual P_{evt} values for the measured quantities $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N$ of the selected events, and the signal fraction f is determined by maximizing \mathcal{L} at that grid point. The likelihood function $\mathcal{L}(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N; m_t, k_{\text{JES}})$ is then projected onto the m_t and k_{JES} axes by integrating using Simpson's rule [22] over k_{JES} and m_t , respectively. Best unbiased estimates of m_t and k_{JES} and their statistical uncertainties are extracted from the mean and standard deviation (SD) of $\mathcal{L}(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N; m_t)$ and $\mathcal{L}(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N; k_{\text{JES}})$.

Further details on the previous implementation of the ME technique can be found in Ref. [10].

3. Reducing the computation demand of the matrix element technique with low-discrepancy sequences in MC integration

The expected uncertainty of the MC integration method based on classical pseudo-random number sequences, decreases as

$$\frac{1}{\sqrt{N}}, \text{ for } N \rightarrow \infty,$$

where N is the number of integral samplings, i.e., points in the $[0, 1]^d$ unit hypercube of dimension d for which the integrand is evaluated [20]. By contrast, LDS converge as $\log^{d-1}(N)/N$ [23], which results in more rapid decrease according to

$$\frac{1}{N}, \text{ for } N \rightarrow \infty.$$

This superior convergence rate is achieved by utilizing a sequence of points that *per constructionem* sample the unit hypercube as uniformly as possible. Thus, LDS are fully deterministic and not random, despite that they are often referred to as “quasi-random numbers”. LDS should not be confused with an arrangement of equidistant points on a lattice, which shows a slower convergence rate for typical applications: for example, $N = n^d$ samplings of the integral are needed to fill a lattice with n points per dimension.

The uniformity of the coverage of the unit hypercube can be quantified rigorously by introducing the mathematical concept of discrepancy D^* . For the purposes of this document, a lower D^* value results in a more uniform coverage of the unit hypercube, and thereby a faster convergence of the MC integration. A

rigorous definition and discussion of the D^* concept is beyond the scope of this document, and can be found in Ref. [23].

The simplest LDS is given by the van der Corput series in one dimension [23], which achieves a uniform coverage of the interval $[0, 1]$ through a consecutive placement of sampling points at $0, \frac{1}{2}, \frac{1}{4}, \frac{3}{4}, \frac{1}{8}, \frac{5}{8}, \frac{3}{8}, \frac{7}{8}$, etc. Among the best performing multidimensional LDS are those given by Faure [24], Sobol [25], and Niederreiter [26], which are all based on the van der Corput series. Based on the findings in Ref. [23], we disregard the Faure sequence. For reducing the computational demand of the ME technique, we tried computer program implementations of the Sobol [27] and Niederreiter [28] sequences provided by the Intel FORTRAN compiler [29]. Both indicate a similar performance in the convergence rate of the MC integration. Most of our findings presented below apply therefore to both the Sobol and Niederreiter sequences. However, the time for the generation of the Sobol sequence is considerably less than for the Niederreiter sequence, and we therefore use the Sobol sequence for our implementation of the ME technique, and as the LDS of reference in this document.

One of the central points in numerical integration is to determine reliably the level of achieved precision: an overly optimistic estimate may result in worsened performance of the method because of its greater numerical uncertainty. However, a too pessimistic estimate will waste computing resources. For the numerical evaluation of the integral G of a function g defined on the unit hypercube $[0, 1]^d$ using MC integration based on pseudo-random numbers, the standard error estimate is often used:

$$\hat{\epsilon}_{\text{std}} \equiv \frac{1}{\sqrt{N}} \left\{ \frac{1}{N-1} \sum_{i=1}^N (g(\xi_i) - \langle g \rangle)^2 \right\}^{\frac{1}{2}}, \quad (3)$$

where ξ_i within $[0, 1]^d$ are the sampling points, and $\langle g \rangle \equiv \frac{1}{N} \sum_{i=1}^N g(\xi_i)$. An alternative, rarely used approach, is to split the original sequence of sampling points ξ_i into K sub-sequences with $\frac{N}{K}$ sampling points, and make K independent integral estimates $G_k \equiv \frac{K}{N} \sum_{i=1}^{N/K} g(\xi_{k+(i-1)K})$, $k = 1, \dots, K$. The error estimate is then given by the sample variance of G_K , i.e.,

$$\hat{\epsilon}_K \equiv \left\{ \frac{1}{K-1} \sum_{k=1}^K (G_k - \langle g \rangle)^2 \right\}^{\frac{1}{2}}, \quad (4)$$

where we have chosen N such that $\frac{N}{K}$ is an integer. The sum $\sum_{k=1}^K G_k$ follows the Student t -statistic [30], which approaches the normal distribution in the limit $K \rightarrow \infty$. Consequently, for finite K , the interval $[\langle g \rangle - \hat{\epsilon}_K, \langle g \rangle + \hat{\epsilon}_K]$ corresponds to a somewhat smaller confidence level than for the normal distribution.

The error estimator $\hat{\epsilon}_{\text{std}}$ in Eq. (3) is not appropriate for LDS, as it is too pessimistic. This follows because Eq. (3) applies to D^* values that are characteristic of pseudo-random numbers, while much smaller D^* are characteristic of LDS. By contrast, the error estimator $\hat{\epsilon}_K$ in Eq. (4) applies also to LDS, under the condition that each of the K sub-sequences used to obtain independent integral estimates G_k is characterised by the same D^* value as the initial sequence.

However, constructing K independent LDS, with same characteristic D^* values is not trivial: for example, randomly assigning each point of the initial sequence to one of the K sub-sequences results in sub-sequences with characteristic D^* values that are different from that of the initial sequence, and from each other. Several involved and sophisticated approaches have been developed to construct subsequences with the same characteristic D^* value as the initial sequence. For instance, the ME technique implemented as described in Ref. [31] uses scrambling [32]. For our implementation of the ME technique, we use the ingeniously simple prescription by Warnock [33]. It utilizes the fact that a LDS of points in $[0, 1]^{d-K}$ can be regarded as K sub-sequences in $[0, 1]^d$, which have the same D^* value *per constructionem*. For our implementation, with $d = 10$, we generate one LDS of $d = 40$, i.e., $K = 4$, which offers a reliable error estimate at a confidence level of about 63%.

Before implementing the LDS in the ME technique, we evaluate their performance and the applicability of error estimators using toy MC integrations of multidimensional test functions: the normal distribution in up to ten dimensions, and a trigonometric function inside a torus of three dimensions. In particular, we define the integrand as $g(\rho) \equiv 1 + \cos(\pi\rho^2/R^2)$, for $\rho \leq r$, and otherwise $g(\rho) = 0$, where ρ is the distance from a given sampling point to the center of the torus tube, with $R = 0.6$ being the distance from the center of the tube to the center of the torus, and $r = 0.3$ the radius of the tube. The integration volume is $V \equiv [-1, 1]^3$. In the following, we focus on the numerically more challenging example of the trigonometric function inside a three-dimensional torus. As a figure of merit, we use

$$\text{relative convergence} \equiv \frac{|\langle g \rangle \cdot V - \int_V f dV|}{\int_V f dV}, \quad (5)$$

where V is the integration volume, and $\int_V f dV = 2\pi^2 R r^2$ is the analytic result.

The relative convergence is compared for the Sobol LDS and the Mersenne-Twister [34] pseudo-random number sequence in Fig. 1 for N up to 1.3×10^8 . As anticipated, the Sobol sequence displays superior convergence behaviour that follows $\frac{1}{N}$. Taking $N = 2^{26} \approx 6.7 \times 10^7$ as an example, the Sobol sequence outperforms the Mersenne-Twister sequence by more than 3 orders of magnitude, and achieves a relative convergence of 5.7×10^{-7} compared to 1.0×10^{-3} .

The performance of the error estimates $\hat{\epsilon}_K$ from Eq. (4) for $K = 4$ and $\hat{\epsilon}_{\text{std}}$ from Eq. (3) using the trigonometric function inside a three-dimensional torus as a test function and the Sobol LDS is evaluated in Fig. 2. Evidently, $\hat{\epsilon}_{\text{std}}$ yields a too pessimistic error estimate for the Sobol sequence, despite that it is applicable to the Mersenne-Twister sequence, as can be seen from comparison with Fig. 1. By contrast, $\hat{\epsilon}_K$ provides an appropriate error estimate for the Sobol sequence, and can therefore be used in our implementation of the ME technique. A practical feature of the $\hat{\epsilon}_K$ estimator is that it dynamically follows the relative convergence, i.e., $\hat{\epsilon}_K$ tends to be small for small values of the relative convergence. This is not the case for $\hat{\epsilon}_{\text{std}}$, which merely gives a monotonously falling upper bound. Thus, $\hat{\epsilon}_K$ can provide a dynamic indication of the achieved nu-




Figure 1: The relative convergence for numerical evaluation of the integral of a trigonometric function inside a torus in three dimensions. The results are obtained using the MC integration technique based on the Sobol LDS and the Mersenne-Twister pseudo-random number sequence.




Figure 2: The relative convergence for numerical evaluation of the integral of a trigonometric function inside a torus in three dimensions. The results are obtained using the MC integration technique based on the Sobol LDS. Also shown are the error estimates $\hat{\epsilon}_K$ from Eq. (4) for $K = 4$ and $\hat{\epsilon}_{\text{std}}$ from Eq. (3).

merical precision through the dips observed in the relative convergence. This feature of $\hat{\epsilon}_K$ is illustrated for $K = 4$ in Fig. 3, for a sub-range in N . We remark that the dips in relative convergence and, consequently, $\hat{\epsilon}_K$ tend to occur for $N = 2^n$, where n is an integer. This is because the unit hypercube is sampled most uniformly for such N . We profit from this feature in various places of our implementation of the ME technique, for example, when we perform the pre-integration step (described in Sec. 2) with $N = 2^{14}$ samplings of the integral.

Having tested the performance of LDS and verified the applicability of the error estimator $\hat{\epsilon}_{K=4}$, we proceed to implement the Sobol sequence in our ME technique. As described in Sec. 2, we use prior information for most of the integration




Figure 3: Same as Fig. 2, however for a sub-range of N , and without showing $\hat{\epsilon}_{\text{std}}$. Powers of 2 are indicated by the thin vertical broken lines in green.

variables by using sampling points distributed according to the prior function, a technique that is commonly referred to as importance sampling. Two common approaches to achieve this are the accept-reject method, and the cumulative distribution function (CDF) method based on the CDF of the prior [20]. It is important to note that the former method, previously employed in Ref. [10], cannot be used in conjunction with LDS. This is because LDS loses its superior property of a low D^* value if a subset of points of the sequence is rejected. By contrast, in the CDF method, all points are preserved and only their mapping from the unit hypercube to the sampling space is modified. No other major changes are required for using the Sobol sequence in our implementation of the ME technique besides switching to the CDF method for sampling the integration space according to a given prior.

After implementing the Sobol sequence for MC integration, we find a reduction of the computation time for the calculation of P_{sig} from about 2 hours per event, averaged over the sample of simulated $t\bar{t}$ events for $m_t = 172.5$ GeV, to about 15 min/event, i.e., by about one order of magnitude. This improvement is for a required numerical precision of 1%, which is found to be sufficient for a robust statistical performance of our implementation. However, our tests with MC integration indicate that the relative gain in computation time may be even greater for smaller required precision.

4. Reducing the computation demand by factoring out the k_{JES} dependence from the matrix element calculation

As already mentioned in Sec. 2, we construct the likelihood on a grid in (m_t, k_{JES}) with spacings of (1 GeV, 0.01) for the extraction of m_t and k_{JES} . For standard samples of simulated MC events which account for a major fraction of the computational demand, this is done for m_t within $[m_t^{\text{gen}} - 12 \text{ GeV}, m_t^{\text{gen}} + 12 \text{ GeV}]$ and for k_{JES} within $[k_{\text{JES}}^{\text{gen}} - 0.1, k_{\text{JES}}^{\text{gen}} + 0.1]$, where m_t^{gen} is the generated m_t and $k_{\text{JES}}^{\text{gen}}$ the generated k_{JES} . Thus, P_{sig} has to be calculated for $25 \times 21 = 525$ grid points in (m_t, k_{JES}) .

In our previous implementation of the ME technique, we recalculated P_{sig} entirely for each point in (m_t, k_{JES}) . However, the integrand in Eq. (2) depends on k_{JES} only via the transfer function $W(\vec{x}, \vec{y}, k_{\text{JES}})$. Furthermore, as detailed in Sec. 2, the integration in Eq. (2) is performed over nine partonic variables and κ , none of which depend on k_{JES} . Therefore, in our new implementation, we factor out the k_{JES} dependence from the ME computation and perform the calculation of $\mu(\vec{y}; m_t) \equiv \int d\sigma(\vec{y}; m_t) d\vec{q}_1 d\vec{q}_2 f(\vec{q}_1) f(\vec{q}_2)$ in Eq. (2) only once for a given sampling point. We then obtain all the 21 integrand values in Eq. (2) for the different k_{JES}^i , $i = 1, 2, \dots, 21$ by multiplying $\mu(\vec{y}; m_t)$ with the transfer function $W(\vec{x}, \vec{y}, k_{\text{JES}}^i)$. Thus, we obtain 21 simultaneous estimates for P_{sig} .

After factoring out the k_{JES} dependence from the ME calculation as described above, we find a further reduction of the computation time for the calculation of P_{sig} from about 15 min/event, after the implementation of LDS, to about 80 s/event, i.e., by another order of magnitude. We note that the reduction is somewhat smaller than the factor of 21 that would be naively expected from the number of grid points in k_{JES} . This

is because of the increased overhead of keeping track of the 21 simultaneous P_{sig} estimates.

The computation time for P_{bkg} is much less of an issue compared to P_{sig} . This is because P_{bkg} does not depend on m_t by definition, and has to be calculated only for 21 points in k_{JES} . Therefore, we did not apply the new approach of factoring out the k_{JES} dependence from the matrix element calculation in P_{bkg} .

5. Validation of the new implementation of the matrix element technique

To verify that the sensitivity of our implementation of the ME technique was not adversely affected by the modifications described in Secs. 3 and 4, we study the response of the ME technique in m_t and k_{JES} . This is done by comparing the extracted m_t^{fit} with the generated m_t^{gen} using pseudo-experiments, and using analogous procedures for k_{JES} . The pseudo-experiments are comprised of $t\bar{t}$ events and dominant background contributions according to their respective fractions measured in data [11]. To evaluate the method's response in m_t , we use five simulated MC samples for $t\bar{t}$ production with $m_t^{\text{gen}} = 165, 170, 172.5, 175$, and 180 GeV for $k_{\text{JES}}^{\text{gen}} = 1$. Similarly, for k_{JES} we use signal and background MC samples with $k_{\text{JES}}^{\text{gen}} = 0.95, 1, 1.05$, and $t\bar{t}$ signal is generated for $m_t^{\text{gen}} = 172.5$. In this validation, we study a representative set of simulated samples used to model data corresponding to 3 fb^{-1} of integrated luminosity.

The response of our implementation of the ME technique in m_t is presented in Fig. 4, before and after the improvements described in Secs. 3 and 4. The results are shown split into $e + \text{jets}$ and $\mu + \text{jets}$ channels defined by the presence of one isolated electron or muon with $p_T > 20$ GeV. Given that the ME technique calculates event probabilities *ab initio* and relies on an analytic parametrization of detector response, its performance is remarkably close to ideal, defined by an offset parameter of zero and by a slope parameter of unity. Most important, the response of the ME technique before and after the improvements is consistent within statistical uncertainties.

As an additional cross-check, we verify that the statistical sensitivity of the ME technique remains consistent. For this, we study the width of the distribution in the pull of m_t . The pull of m_t is defined as $\sum_{i=1}^{1000} (m_t^i - \langle m_t \rangle) / \sigma_{m_t}^i$, where m_t^i and $\sigma_{m_t}^i$ are the extracted values of m_t and its corresponding statistical uncertainty found in pseudo-experiment i , $\langle m_t \rangle \equiv \frac{1}{1000} \sum_{i=1}^{1000} m_t^i$, and the sums extend over all 1000 pseudo-experiments conducted for a given m_t^{gen} . As Fig. 5 demonstrates, the statistical sensitivity is within 20% of the ideal pull width of unity, and is consistent within statistical uncertainties before and after the improvements.

In similar spirit, we study the response and statistical sensitivity in k_{JES} before and after the improvements of Secs. 3 and 4. The results for the response are shown in Fig. 6, while the statistical sensitivity is presented in Fig. 7. Both figures display a consistent performance of the ME technique before and after our improvements.

Our validation studies in m_t and k_{JES} indicate full consistency within statistical uncertainties. We therefore conclude that our




Figure 4: The response of the ME technique in m_t obtained using pseudo-experiments constructed from MC events with fully simulated response of the D0 detector. Each data point corresponds to the mean extracted m_t averaged over 1000 pseudo-experiments at a given m_t^{gen} . The dependence is fitted with a linear function (black solid line), with the ideal case indicated as the red broken line. The results obtained with our previous implementation of the ME technique are shown in (a) for the e +jets and in (b) for the μ +jets channel. Analogous results obtained including the improvements described in this manuscript are shown in (c) and (d).




Figure 5: The statistical sensitivity of the ME technique in m_t obtained using pseudo-experiments constructed from MC events with fully simulated response of the D0 detector. Each data point corresponds to the width of the distribution in the pull of m_t found with 1000 pseudo-experiments at a given m_t^{gen} . The dependence is fitted with a constant (black solid line), with the ideal case indicated as the red broken line. The results obtained with our previous implementation of the ME technique are shown in (a) for the e +jets and in (b) for the μ +jets channel. Analogous results obtained including the improvements described in this manuscript are shown in (c) and (d).




Figure 6: Same as Fig. 4, but for k_{JES} .




Figure 7: Same as Fig. 5, but for k_{JES} .

implementation of the ME technique using LDS for the MC integration and factorizing the k_{JES} dependence from the ME calculation does not adversely affect the performance of the method, and can be applied for data analysis.

6. Conclusion

In conclusion, we have presented the numerical integration approaches implemented to reduce by a factor of 90 the computational demand of the calculation of event probabilities using the ME technique. We achieve this by using low-discrepancy sequences for the numerical MC integration in conjunction with a dedicated estimator of the numerical uncertainty — a novelty in the context of the ME technique, as well as the factorization of the jet energy scale factor k_{JES} from the ME calculation — newly applied in the context of m_t measurements with an *in situ* jet energy scale calibration. These improvements have been validated through MC studies. The low-discrepancy sequences are universally applicable for numerical MC integration, and are not specific to the presented studies.

Acknowledgments

We thank our D0 colleagues for useful discussions and for their kind permission to use the D0 detector simulation and other collaborative software to expedite the preparation of this paper. The authors acknowledge the support from the Department of Energy (USA), the National Science Foundation (USA), the Bundesministerium für Bildung und Forschung (Germany), and the Deutsche Forschungsgemeinschaft (Germany).

References

- [1] K. Kondo, J. Phys. Soc. Jap. 57 (1988) 4126–4140.
- [2] J. Neyman, E. S. Pearson, Phil. Trans. R. Soc. Lond. A 231 (1933) 289–337.
- [3] V. Abazov, et al., Nature 429 (2004) 638–642.
- [4] V. Abazov, et al., Phys. Lett. B617 (2005) 1–10.
- [5] V. Abazov, et al., Phys. Rev. Lett. 103 (2009) 092001.
- [6] T. Aaltonen, et al., Phys. Rev. Lett. 103 (2009) 092002.
- [7] T. Aaltonen, et al., Phys. Rev. Lett. 103 (2009) 101802.
- [8] G. Aad, et al., Phys. Lett. B 716 (2012) 30 – 61.
- [9] P. Artoisenet, V. Lemaître, F. Maltoni, O. Mattelaer, JHEP 1012 (2010) 068.
- [10] V. M. Abazov, et al., Phys. Rev. D 84 (2011) 032004.
- [11] V. M. Abazov, et al., Phys. Rev. Lett. 113 (2014) 032002.
- [12] V. M. Abazov, et al., Phys. Rev. D 84 (2011) 052005.
- [13] J. Pumplin *et al.*, J. High Energy Phys. 07 (2002) 012.
- [14] T. Sjöstrand *et al.*, Comp. Phys. Commun. 135 (2001) 238.
- [15] T. Sjöstrand, S. Mrenna, and P. Skands, J. High Energy Phys. 05 (2006) 026.
- [16] G. Mahlon, S. J. Parke, Phys. Rev. D53 (1996) 4886–4896.
- [17] G. Mahlon, S. J. Parke, Phys. Lett. B411 (1997) 173–179.
- [18] J. Hammersley, D. Handscomb, Methuen, London (1964).
- [19] F. James, Comput. Phys. Commun. 79 (1994) 111.
- [20] W. Press, et al., Numerical Recipes in Fortran 77: The Art of Scientific Computing, Cambridge University Press, New York, NY, USA, 2 edition, 1992.
- [21] F. A. Berends *et al.*, Nucl. Phys. B 357 (1991) 32.
- [22] K. Atkinson, An Introduction to Numerical Analysis, Wiley, 1989.
- [23] J. Bossert, et al., Nucl. Instrum. Meth. A559 (2006) 232–236.
- [24] H. Faure, Acta Arith. 41 (1982) 337.
- [25] I. Sobol, USSR Comput. Math. Phys. 16 (1967) 236.
- [26] H. Niederreiter, Monatsh. Math. 104 (1987) 273.
- [27] P. Bratley, B. L. Fox, ACM Trans. Math. Softw. 14 (1988) 88–100.
- [28] P. Bratley, B. L. Fox, H. Niederreiter, ACM Trans. Model. Comput. Simul. 2 (1992) 195–213.
- [29] Intel FORTRAN compiler with Math Kernel Library 11.0, <http://software.intel.com/en-us/articles/intel-fortran-composer-xe-documentation/> (2011).
- [30] “Student” [William Sealy Gosset], Biometrika 6 (1908) 1.
- [31] Aaltonen, T. *et al.*, Phys. Rev. Lett. 105 (2010) 252001.
- [32] A. B. Owen, in: Proceedings of the 1998 Winter Simulation Conference, Press, 1998, pp. 571–577.
- [33] T. Warnock, LA-UR-01-1950 (2001).
- [34] M. Matsumoto, T. Nishimura, ACM Trans. Model. Comput. Simul. 8 (1998) 3–30.